

CLAIMS

What is claimed is:

1. A method comprising:
 - maintaining a destination state in a destination of a network, the destination state
 - 5 being in a first state;
 - receiving, in the destination, a first set of packets of a plurality of packets belonging the first state;
 - receiving, in the destination, a barrier indication belonging to a second state;
 - manipulating at least one packet of the first set of packets in response to the
 - 10 received barrier indication.
2. The method of claim 1, wherein manipulating the at least one packet includes sending or output queuing the at least one packet.
3. The method of claim 2, further comprising storing the at least one packet in the
- 15 4. The method of claim 3, wherein each of the first set of packets includes a unique sequence number within the first set of packets; and wherein the first set of packets are sent or forwarded from the destination in an order based on said unique sequence numbers.
5. The method of claim 4, further comprising receiving, in the destination, a
- 20 second set of packets of the plurality of packets belonging to the second state; wherein each of the second set of packets has a unique sequence number within the second set of packets, and wherein the unique sequence number of a packet within the first set of packets is the same as the unique sequence number of a packet within the second set of packets.

6. The method of claim 1, wherein manipulating the at least one packet includes resequencing or reassembling the at least one packet.

7. The method of claim 1, wherein the barrier transition is received in a packet or a signal over a control or data line.

5 8. The method of claim 1, wherein receiving the barrier indication includes receiving and aggregating a plurality of messages or signals received from a plurality of elements of the network.

9. The method of claim 8, wherein the plurality of elements of the network include a plurality of planes of a switching fabric.

10 10. The method of claim 8, wherein the plurality of elements of the network include a plurality of switching elements of a switching fabric.

11. The method of claim 1, further comprising updating the current value of the destination state with the current value of the identifiable state.

15 12. The method of claim 1, further comprising maintaining a source state in a source of the network.

13. The method of claim 1, wherein the plurality of packets include data and control packets.

14. The method of claim 1, wherein the network includes a packet switching fabric

20 15. The method of claim 1, wherein the network includes a packet switching system.

16. A packet switching system performing the method of claim 1.

17. A router performing the method of claim 1.

18. A computer-readable medium containing computer-executable instructions for performing the method of claim 1.

19. A packet switch performing the method of claim 1 for a plurality of
5 independent sets of packets.

20. The packet switch of claim 19, wherein the plurality of independent sets of packets are identifiable based at least on a priority or type of traffic.

21. A method for routing a plurality of packets between a first source of a plurality of sources and a first destination of a plurality of destinations in a packet switching
10 system, the method comprising:

each of the plurality of sources maintaining at least one source state indication;

each of the plurality of destinations maintaining at least one destination state indication;

the first source sending a first set of packets of the plurality of packets to the first
15 destination, each packet of the first set of packets including a first state indication reflected in one of the at least one source state indications maintained in the first source;

after sending the first set of packets, the first source sending a source barrier message over each of a plurality of paths leading to the first destination through the packet switching system;

20 in response to receiving a destination barrier message generated in response to said sent plurality of source barrier messages, the first destination recognizing a barrier transition or a new barrier state and updating one of the at least one destination state indication; and

the destination manipulating at least one packet of the first set of packets in
25 response the barrier transition or the new barrier state.

22. The method of claim 22, further comprising:

each of the plurality of destinations relaying a barrier state indication of a barrier transition or the new barrier state to at least one of the plurality of sources;

in response to receiving a first predetermined number of barrier state indications,
5 each of the at least one of the plurality of sources sending a source barrier acknowledgement message over each of a plurality of paths leading from the at least one of the plurality of sources;

after receiving a second predetermined number of destination acknowledgement messages, each of the plurality of destinations relaying a barrier acknowledgement
10 indication to the at least one of the plurality of sources; and

in response to receiving the first predetermined number of barrier acknowledgement indications, each of the at least one of the plurality of sources updating at least one of said at least one source state indication.

23. A method performed by a switching element of a packet switching system, the
15 method comprising:

maintaining a plurality of output queues for containing packets to be sent over a plurality of outgoing links to a plurality of other components of the packet switching system

maintaining a barrier state;
20 updating the barrier state in response to receiving a barrier state transition request on each of a plurality of incoming links;

placing a packet containing the barrier state transition request in at least each of the output queues that are occupied by at least one packet; and

25 sending packets over the plurality of outgoing links to the plurality of other components of the packet switching system.

24. The method of claim 23, wherein sending packets includes sending a packet containing the barrier state transition request out each of the outgoing plurality of links.

25. The method of claim 23, further comprising placing a packet containing the barrier state transition request in at least each of the output queues that are occupied by at least one packet.

5 26. The method of claim 23, further comprising sending an outgoing barrier acknowledgement message out each of the outgoing plurality of links in response to receiving an incoming barrier acknowledgement message over said each of a plurality of incoming links.

27. A packet switching system comprising:
means for maintaining a barrier state in a plurality of components;
10 means for updating the barrier state in the plurality of components;
means for determining whether or not to send packets indicating a particular barrier state from a particular component of the plurality of components based on a current or previous barrier state maintained by the particular component.

28. The packet switching system of claim 27, wherein said means for updating the
15 barrier state includes means for propagating barrier request messages.

29. The packet switching system of claim 27, wherein said means for updating the barrier state includes means for propagating barrier acknowledgement messages

30. A packet switching system comprising:

a plurality of source elements, each of the plurality of source elements including: a source barrier state maintainer to indicate a current source state; a barrier accumulator to receive indications of a first subset of a plurality of barrier request and acknowledgement messages, to determine when a particular barrier request or acknowledgement message
5 may be sent, and to update the current source state;

a plurality of switching nodes, each of the plurality of switching nodes including: a switching node barrier state maintainer to indicate a current node state; a switching node barrier accumulator to receive indications of a second subset of the plurality of
10 barrier request and acknowledgement messages, to determine when a second particular barrier request or acknowledgement message may be forwarded to a next switching node or destination element, and to update the current node state; and

a plurality of destination elements; each of the plurality of destination elements including: a destination barrier state maintainer to indicate a current destination state; a
15 destination barrier accumulator to receive a third subset of the plurality of barrier request and acknowledgement messages, to determine when a particular barrier request or acknowledgement message may be relayed to at least one of the plurality of source elements, and to update the current destination state; and a packet sending mechanism to determine when a particular packet associated with a particular state may be manipulated
20 or sent from the particular destination element based on the current or a previous destination state.